Using 3DM Analyst Mine Mapping Suite for Slope Stability — Case Studies

J.S. Birch *ADAM Technology, Australia*

Abstract

Using photogrammetry to generate 3D surface models for structural mapping is now a well established technique and several vendors provide off-the-shelf software capable of performing this task. In this paper we focus on the task of acquiring the necessary photographs and processing them to derive the surface model ready for mapping, including situations where speed is of the essence and situations where placing control points for georeferencing is impossible.

1 Introduction

Photogrammetry is the science of using sets of 2D images to reconstruct 3D data. In order to do this, the precise location and direction of each image, the *exterior orientation*, must first be determined.

While measuring the camera's *position* accurately enough is not too difficult, measuring the direction that the camera was pointing in directly is far more challenging; for a 200 mm lens, the accuracy of the angular measurement would need to be about 1 arc-second — the accuracy of the best total stations — to avoid degrading the accuracy of the resulting data, and even then it would be highly susceptible to knocks and vibration, and *still* not accurate enough for longer focal length lenses. It is therefore highly desirable that the software be able to *automatically* determine the direction accurately, simply by analysing the images themselves. It is also quite desirable for it to automatically determine the position as well, since this allows workflows where measuring the camera's position is impractical (for example, photography from a moving platform).

This does not remove the need to somehow georeference the data into the desired co-ordinate system, however. The usual approach is to use *control points* — points in the scene with known co-ordinates, commonly determined using a total station or GPS. Measured camera locations can be used instead of control points if required, although as will be shown later there are good reasons for preferring control points where possible.

One other requirement for the software to determine 3D co-ordinates for all of the points of interest is that those points must be visible in at least two images taken from two different locations. In practice this is not difficult to achieve, and there are many standard camera configurations that can be used to minimise the effort required to capture the images while avoiding the risk of failing to satisfy this requirement.

In this paper we will explore a number of examples with varying camera configurations and georeferencing techniques.

2 Terrestrial photography

Taking photos from the ground is the cheapest and simplest method, and has the advantage that the faces of pit walls are generally close to perpendicular to the view direction and that surveying camera stations is an option. In this section we will look at several different scenarios with different combinations of photographic techniques and control point use.

All software processing times in this section are from a 2.4 GHz Intel Core 2 Duo PC and therefore represent the performance obtainable on a medium- to high-end laptop computer.

2.1 Medium range

A 390 m × 190 m section wall was captured from two camera stations approximately 450 m away on the crest of the opposite wall and 289 m apart (Figure 1). To capture the entire area with a Canon EOS 5D and 100 mm lens required six images from each location, giving a ground pixel size of approximately 3.7 cm, a fairly typical resolution for geotechnical analysis. (ADAM Technology calls this technique “image fanning”, which is best used with longer focal length lenses.)

The total time required to capture both sets of images was just under 25 minutes, including walking to and from each location, setting up the tripod, and packing up again. Most of the time was actually spent walking, so being able to drive to each location (or locating the camera positions closer together) would reduce the image capturing time significantly.

Prior to capturing the images, six control points were painted on the wall — three near the top and three at the bottom. (Three control points in total — or, if preferred, one control point plus both camera stations — would have been sufficient, but would lack redundancy; having six allows the software to detect any errors in the control point data or in indentifying the control points on the images.) To make identification simpler, the numeric ID of each control point was painted on the wall next to it. They were surveyed by a reflectorless total station located roughly halfway between the two camera stations. Although placing the control points took some time (mostly because of the need to walk approximately 700 m each way along the upper bench to paint them), experience shows that they should remain in place for quite a while, so this time can be amortised over many capture sessions. (Walking the pit painting control points is also a good opportunity to assess the rock at close range, gaining insight that can be used later, when mapping.)

A number of prisms were also mounted on the wall for monitoring purposes; these can be used as control points, too, if they are visible in the images. In this particular case, a longer focal length lens would be required to achieve a pixel size small enough to show the prisms. Mounting a metal plate around the prisms to make it easier to see them and allow them to be used in jobs where a shorter focal length is sufficient.

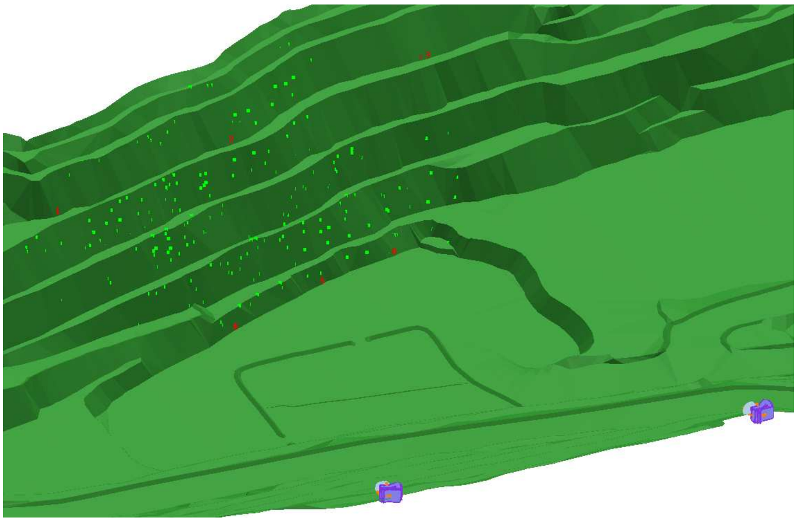

Figure 1 Screenshot of 3DM CalibCam showing the area captured, with the pit design model for context. Control points are in red. Light green points are points automatically detected by the software to determine the camera orientations with respect to each other. These are known as “relative-only” points, because their 3D locations are unknown beforehand. Without the control points the result would look exactly the same except the co-ordinate system would be arbitrary.

Note that control points don’t actually need to be painted; anything in the images for which 3D co-ordinates are known can be used as control. As will be shown in a later example, a reflectorless total station and easily identifiable features can eliminate the need to approach the wall at all, and control points can also be both in front of the wall (e.g. witches’ hats on the ground) and behind the wall (e.g. monuments around the outside of the pit).

Camera positions were not surveyed. One advantage of relying entirely on control points rather than surveying camera positions in lieu of control points is that images can be captured again in the future without needing to set up over the same locations and without requiring additional surveying. It also makes it easier for the software to compensate for the most common types of camera calibration errors, which will be discussed in more detail later.

Images were captured in JPEG format; this eliminates any processing step between image capture and loading the images into 3DM Analyst Mine Mapping Suite, increasing productivity and eliminating a potential source of errors. Although theoretically JPEG can have an impact on the data generated, in practice the effect is negligible if the high quality setting is used on the camera, and the workflow advantages from using JPEG are substantial.

Figure 2 Surface model generated by 3DM Analyst. The six control points are in green.

From the time of the first image to the last image was 28 seconds for the left camera station and 22 seconds for the right. Image capture can be quick because it is not necessary to try and line up the images from the second camera station with the corresponding images from the first — all that is required is to ensure the images captured from the same location overlap each other. The auto-focus marks in the viewfinder are useful for this: simply locate a point on the surface that is under the right-most auto-focus mark and then rotate the camera until that same point is under the left-most auto-focus mark. This gives an overlap of about 30%. (In this project, a point roughly halfway between the edge of the frame and the most extreme auto-focus mark was chosen to give overlaps of about 20% both horizontally and vertically.)

The reason why it is not necessary to line up corresponding images is because 3DM CalibCam, the program in 3DM Analyst Mine Mapping Suite responsible for determining exterior orientations, allows images captured from the same location to be merged together into one or more higher-resolution images, similar to panoramic software that is available with many digital cameras (although, in this case, the merging is photogrammetrically correct to ensure there is no reduction in accuracy). There is a trade-off in doing this — larger images take longer to process and place a greater burden on the computer’s graphics, processor, and memory subsystems, so the user can choose an image size to work with that fits the power of the computer they are using.

In this particular project we decided to process it twice: once using a single merged image of approximately 60 megapixels for each camera station, and once with the original images from one camera station and a single merged image for the other. (By matching original images with a single merged image we eliminate the problem of misalignment of individual images, while still keeping each DTM (Digital Terrain Model, also known as a “3D Image”) small, because each DTM only covers the area covered by *both* images, and any area in one image that is outside the other is automatically ignored.)

Processing times (in minutes:seconds) on the 2.4 GHz dual-core PC were as follows:

1. 3DM CalibCam Project Setup: 0:30

2. Automatic RO Point Generation*, Relative Orientations: 2:15
3. Control Point Digitising, Absolute Orientations: 3:00
4. Merged Image Generation*: 0:55
5. 3DM Analyst Project Generation: 0:15
6. DTM Generation* (Figure 3): 7:10

Total Time: 14 minutes 5 seconds (User Time: 4 minutes.)

Total Points Generated: 1,804,120.

(Items with an asterisk above and in subsequent lists are those that are automated and which will speed up if a faster PC is used. Times include time taken to analyse the reports and verify that the results were acceptable. The images loaded into the software in step 1 were the images as captured by the camera; there was no processing step between image capture and step 1.)

It is worth noting that the DTMs that are generated aren't simply raw point clouds, as a laser scanner might produce — the above time *includes* triangulating the point cloud and filtering out any bad points, which actually takes up a fair bit of time. They are also already in the desired co-ordinate system, so no further processing is required — the projects are ready to be loaded into 3DM Analyst for mapping.

For the two merged images case, the first three steps were the same but the remaining steps were:

4. Merged Image Generation*: 1:45
5. 3DM Analyst Project Creation: 0:35
6. DTM Generation*: 8:30

Total Time: 16 minutes 35 seconds (User Time: 4 minutes 20 seconds.)

Total Points Generated: 1,169,198.

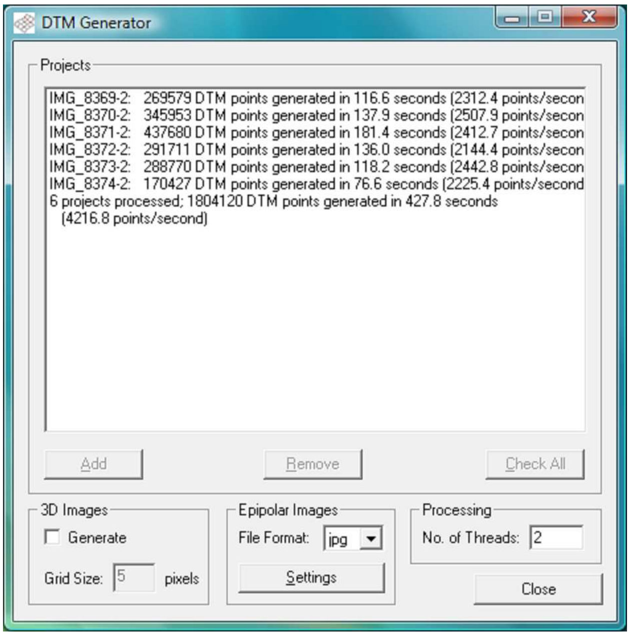

```
IMG_8369-2:   269579 DTM points generated in 116.6 seconds (2312.4 points/secon
IMG_8370-2:   345953 DTM points generated in 137.9 seconds (2507.9 points/secon
IMG_8371-2:   437680 DTM points generated in 181.4 seconds (2412.7 points/secon
IMG_8372-2:   291711 DTM points generated in 136.0 seconds (2144.4 points/secon
IMG_8373-2:   288770 DTM points generated in 118.2 seconds (2442.8 points/secon
IMG_8374-2:   170427 DTM points generated in 76.6 seconds (2225.4 points/second
6 projects processed; 1804120 DTM points generated in 427.8 seconds
  (4216.8 points/second)
```

3D Images: ☐ Generate; Grid Size: 5 pixels

Epipolar Images: File Format: jpg

Processing: No. of Threads: 2

Figure 3 Batch processing the DTMs for the first case. Note that because a dual-core PC was used, the projects were processed two at a time. DTM Generator can process up to eight projects at once if the computer has sufficient memory.

Note that the total number of points generated in the two projects is different; this is largely due to the overlaps present in the first case (because the original images were overlapped by 20% on each side, so, too, are the DTMs) that are eliminated in the second by merging the images before generating the DTM. The speedup that this brings (avoiding processing the same area multiple times) is offset by the fact that some operations during DTM generation cannot be spread across multiple CPUs, so using two processors to process one DTM has a lower throughput than using two processors to process two different DTMs. Also, the single image DTM has more points in it, and adding an extra point to a DTM that has over a million points in it already takes longer than adding an extra point to a DTM with only 200,000 points.

Table 1 and Table 2 contain important information derived from the exterior orientation report generated by 3DM CalibCam. The control point residuals are the differences between the derived ("Adjusted") 3D co-ordinates of the control points and the original, supplied 3D co-ordinates. When determining the exterior orientation of each image, the software finds the best fit (in a "least squares" sense) to the control points, the image co-ordinates of the control points, and the relative-only points generated by the software, taking into account the specified accuracy of each co-ordinate of each control point and the specified image co-ordinate accuracy (0.25 pixels in this case). These residuals give us an indication of the accuracy of the resulting data, subject to the constraint that you can't determine the accuracy of something to a higher degree than the accuracy of the tool you are measuring it with — in this case, the control points, which have a stated accuracy of 5 cm. (In other words, if the control points really are accurate to 5 cm, then we would expect a perfect measuring system with no errors to report residuals of 5 cm.)

Table 1 Control point residuals

Control Point IDs	Adjusted Data (m)			Residuals (m)		
	X	Y	Z	X	Y	Z
1	3715.729	4947.612	741.747	0.060	-0.005	0.033
2	3561.993	4920.177	742.300	0.006	-0.004	-0.009
3	3384.305	4884.536	741.457	0.025	0.035	-0.035
4	3599.248	5013.700	658.142	-0.053	-0.019	-0.023
5	3515.392	4989.175	657.289	-0.040	0.001	0.017
6	3451.717	4978.425	657.804	0.001	-0.007	0.016
	Control Point RMS			0.042	0.018	0.026
	Total RMS				0.053	

We would expect a 3D accuracy from a project of this nature of about 2 cm in plan (east-west and up-down in this case, i.e. "X" and "Z") and about 3 cm in depth (north-south in this case, i.e. "Y"), so the figures in Table 1 are reasonable, given a control point accuracy of 5 cm and the ability with which the centre of a large cross painted on a wall can be digitised in an image and also located by the surveyor with the total station. (This last point is worth keeping in mind; the total station may be able to achieve an accuracy of 2 cm, but if the surveyor can't find the centre of the cross accurately then the accuracy of the survey data will suffer. An advantage of using circular targets with fine cross-hairs painted in their centre is that not only can they be easily and accurately digitised in the images, but the surveyors can locate their centres more accurately as well.)

Table 2 Derived camera locations

Camera Station	Camera Location (m)					
	X	Y	Z	σ_X	σ_Y	σ_Z
Left	3693.841	5358.047	797.041	0.096	0.088	0.155
Right	3404.893	5370.095	797.858	0.108	0.088	0.177

Table 2 shows the camera locations that the software derived, and the accuracy with which the software thinks it derived them. The most important thing to notice about the derived camera locations, other than their location being reasonable[1], is the amount of certainty that the software is reporting in their location, i.e. the sigma values. In this case, the software is reporting a height accuracy in the order of 15–20 cm; at a distance of 450 m, this translates to a possible distance error at the top and bottom of the wall of about 4 cm, and a possible dip error for geotechnical analysis of about 0.02° — about two orders of magnitude more accurate than traditional techniques.

2.2 Close range, camera in moving vehicle

This trial was a test of the software's ability to handle images captured from a moving vehicle, driving up the ramp. The vehicle was driven continuously over a distance of 243 m while 20 images were captured out the window by one of the onsite geotechnical engineers using a Canon EOS 5D with a 28 mm lens. The time between capturing the second image[2] and the last image was 1 minute 36 seconds, or about one image every five seconds. To determine correct spacing, the photographer simply observed the wall through the viewfinder and triggered the camera when a point near the left edge got close to the centre of the image.

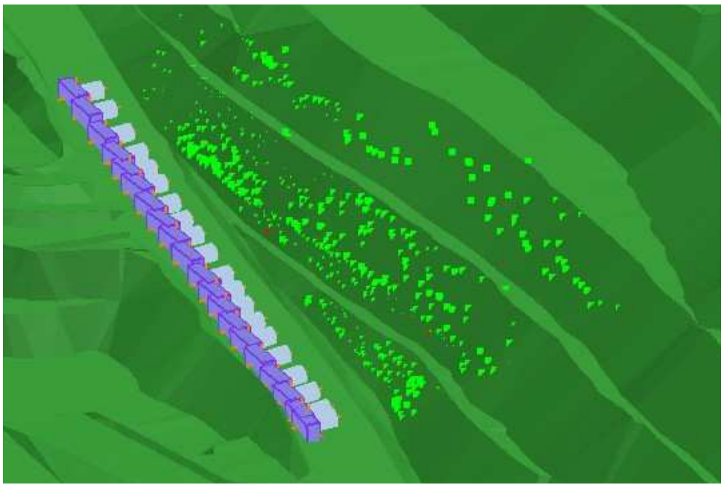

Figure 4 Camera orientations for the images captured out the window of a moving vehicle.

The key requirement for capturing images from a moving vehicle is to keep the shutter speed high to avoid motion blur[3]. In this case the images were captured at 1/250th of a second, which was sufficient. If it was overcast, or late in the afternoon, a larger aperture or higher ISO setting might be required to achieve this.

Processing times were as follows:

1.	3DM CalibCam Project Setup:	0:20
2.	Automatic RO Point Generation*, Relative Orientations:	7:10
3.	Control Point Digitising, Absolute Orientations:	2:55
4.	3DM Analyst Project Generation:	1:00
5.	DTM Generation*:	8:15

---

[1] That is, they're not impossible; they don't need to be *correct*, because one of the largest benefits of *not* surveying the camera locations is the ability this gives the software to compensate for certain types of calibration error by locating the cameras either closer to or further away from the wall than they really were, as we'll discuss later.

[2] The time between the first and second images was longer as the technique was being explained.

[3] The very expensive large-format cameras traditionally used for aerial photography have a feature known as "forward motion compensation", or FMC. This literally shoots the camera backwards at the same speed the aircraft is flying at during the exposure to eliminate motion blur. Lacking that feature the only alternative is to ensure a sufficiently high shutter speed.

Total Time: 19 minutes 40 seconds (User Time: 8 minutes 15 seconds.)

Total Points Generated: 1,981,549.

Since there were only three control points visible in this trial, the residuals in this case are not useful. (They are not zero, because there *is* a very small amount of redundancy with three control points, but it is so small that it has no value for estimating project accuracy.)

Figure 5 Image captured out the open window of a moving vehicle. Note the control point painted on the wall, and the vehicle's window frame on the right.

Figure 6 and Figure 7 demonstrate that these images can be used for structural mapping.

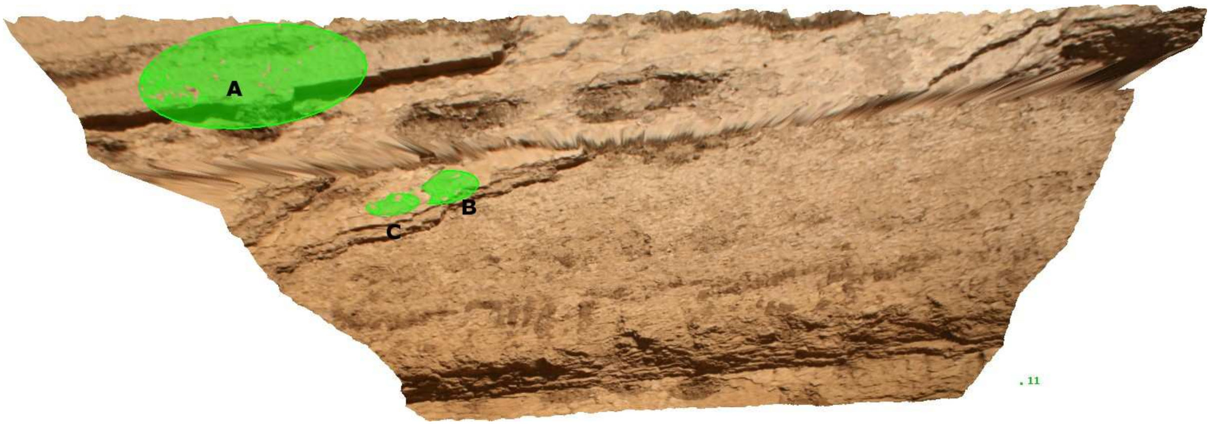

Figure 6 Features digitised on the wall above the ramp.

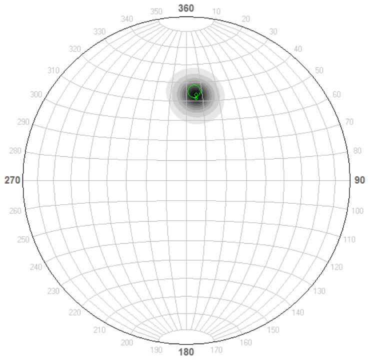

Figure 7 3DM Analyst stereonet of the features in Figure 6.

Table 3 Feature details

Feature	X	Y	Z	Dip	Dip Dir'n	Max. Chord
A	3103.887	5291.144	726.260	45.4	185.2	11.521
B	3111.477	5274.612	713.921	43.7	186.9	3.814
C	3108.898	5273.637	712.866	41.5	187.0	2.706

2.3 Medium range, no marked control points

For this project, a 400 m × 280 m section of pit wall was captured from about 600 m away using two camera positions 93 m apart. Images were captured with a 50 mm lens (two images from each station), a 100 mm lens (six images from one station and seven from the other), and a 200 mm lens (21 images from one station and 27 images from the other), yielding a ground pixel size of approximately 10 cm, 5 cm, and 2.5 cm, respectively.

No control points were marked on the wall at all; instead, the two camera stations plus a delineator on top of the wall were surveyed using GPS. One possibility that was not tested is the use of a flash to highlight prisms on the wall so they can be used as control points. It is not certain that the reflectors would show up at that range but Fugro have used retro-reflective targets with the flash enabled in daylight when monitoring William Street in the Perth CBD to good effect.

It took 5 seconds at each camera station to capture the 50 mm images, 25 seconds to capture the 100 mm images, and 1 minute 45 seconds to capture the 200 mm images. Changing the lens took 1 minute 18 seconds on average. The time from the last image at the first camera station to the first image at the second camera station was 4 minutes exactly, including the time to pick up the first camera station using GPS. Capturing all three sets of images took about 15 minutes, including surveying the camera stations, and would have taken less than 10 minutes if the 200 mm images were not also captured.

Table 4 shows the image processing times for the various projects, simplified a little — steps 1–4 in the previous list are shown as “Orientations”, with the remaining steps shown as “DTM Generation”. The 100 mm project was processed twice, once with one merged image matched against all of the original images from the other camera station, the other time with two merged images. As before, the latter case reduces duplication (the image merging eliminates the overlap between images) but results in larger data sets, requiring a more powerful PC to handle them. The other projects were only processed with one merged image matched against the original images from the other camera station because it’s faster and easier that way.

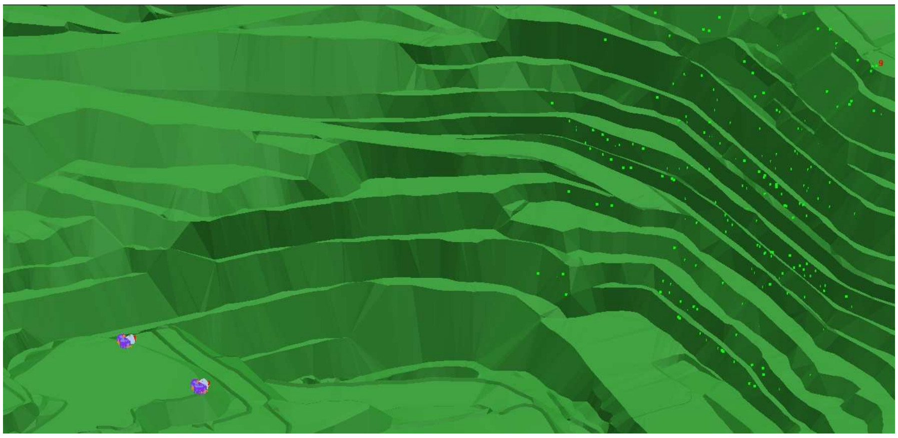

Figure 8 Project location. The red control point on top of the wall represents the delineator that was surveyed, along with the two camera locations.

Table 4. Image Processing Times. "Orientations" includes merged image creation and project generation.

	50 mm (1 merged image)	100 mm (1 merged image)	100 mm (2 merged images)	200 mm (1 merged image)
Number of Images	4	13	13	48
Orientations (m:s)	3:00	4:45	5:25	16:15
DTM generation (m:s)	2:37	7:00	8:30	33:30
Number of Points	764,115	2,512,937	1,651,137	7,394,674

With the project's configuration and good control points we would expect the overall 3D point accuracy of the 50 mm, 100 mm, and 200 mm projects to be 230 mm, 115 mm, and 60 mm, respectively. The biggest factor in the accuracy in this case is the camera separation, which was relatively low. Simply doubling the separation would increase the overall accuracy to 120 mm, 60 mm, and 30 mm, respectively. Although there are often constraints on where images can be captured from, with the relatively long ranges that can be used with photogrammetry it is generally possible to find vantage points that deliver the desired geometry.

In this particular case the absolute accuracy is likely to be lower, partly due to the GPS accuracy, and partly because camera stations were used in lieu of control points. As mentioned earlier, there are advantages to using control points rather than surveying camera stations:

1. When control points are placed near the area to be mapped, survey errors are not magnified by extrapolation.
2. Some types of error in the camera's calibration (in particular the focal length, caused by not physically locking the lens onto a single focal length) can be compensated for by the software if camera positions are not surveyed.
3. Control points tend to be more permanent; if only existing control points are sufficient for a project, then the fieldwork consists of nothing more than capturing some images. If camera stations are being surveyed then the fieldwork becomes more onerous.

2.4 Close range

This trial consisted of very close range images (the base of the wall was just 85 m away) and two different approaches to georeferencing the data — one using three camera positions and no control points at all, the other using existing features picked up by a reflectorless total station located in the same area as the camera. The purpose of this trial was to demonstrate that the software can be used in instances where there are no control points available and no way to place any in the scene — in this case, because of a failure. The trial is somewhat artificial because the ideal solution would have been to simply photograph the failure from further away (e.g. the top of the opposite wall, like in earlier trials) and then control points could have been placed freely in suitable locations.

2.4.1 No control points

Three images were captured with the 28 mm lens in a triangular configuration with the GPS receiver held above the camera while the images were being captured (Figure 9). (The camera was hand-held; no tripod was used.)

Capturing the images took 1 minute 54 seconds from the first image to the last and processing them took 2 minutes 15 seconds (1 minute for orientations, 1 minute 15 seconds for DTM generation) to generate a DTM consisting of 254,171 points.

The biggest advantages of this approach are the speed (both in the field and in processing the images) and the ability to orient all of the data in the desired co-ordinate system without any placement or surveying of control points at all. Using GPS to survey the camera positions also avoids setting up a total station, which took about 30 minutes in this case due to some difficulty backsighting.

The biggest disadvantage is the relative inaccuracy of the data compared to the traditional approach of using control points. If more accuracy is required then there are three options:

1. Fix the focus setting on the lens; this should always be done in cases where camera stations are surveyed since the ability of the software to compensate for small focal length errors is eliminated.
2. Increase the separation between cameras, changing lens if required to make the ground pixel size approximately the same.
3. Increase the number of camera locations.

The second point is worth expanding on because normally ADAM recommends limiting the difference in ground pixel size between images captured from different locations to a factor of no more than 2. This would seem to limit the maximum difference in distances between camera locations to a factor of 2 as well, however that would only be the case if a single lens was used. Since the *ground* pixel size depends not only on the distance from the camera but also the focal length of the lens, a camera position twice as far away from the wall will still have the same ground pixel size if a lens with twice the focal length is used! This means that by using a wide range of camera distances and many camera locations it is possible to achieve any accuracy required at the wall, even if relatively inaccurate surveying methods are used for the camera stations. (The lower the surveying accuracy, the greater the number of camera positions required to achieve a given level of accuracy at the wall; the improvement is proportional to the square root of the number of surveyed locations, all other things being equal.)

2.4.2 Natural points as control

Ten points were picked up by the total station to act as control points. Although these weren’t marked, their locations were generally obvious thanks to the 28 mm project — once an absolute orientation had been obtained using the camera positions, the software’s driveback tool could be used to identify the feature on the pit wall or other object (including one signpost and one delineator) that had been surveyed. This feature could then be digitised in the images from the other projects to georeference them, even though their camera positions were not surveyed. (If there was no other project to fall back on then the surveyor would have to remember at least three of the locations that were picked up by the total station so they could be identified again in the images later. Of course, if the software was being used in the field then the surveyor could identify each point as they were being picked up.)

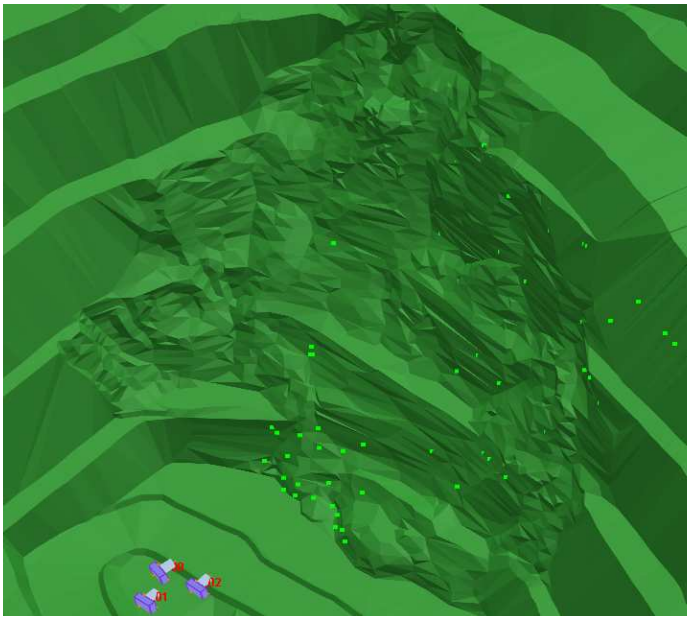

Figure 9 Using camera stations for georeferencing. Increasing the number of camera stations and the distance between them will increase the accuracy of the data at the wall.

The wall was captured three times (in addition to the 28 mm project above) using a 50 mm lens (5 images), a 100 mm lens (17 images), and a 200 mm lens (70 images). (More images were required this time because the field of view was wider due to the proximity of the wall. Normally we would prefer to stand further back, especially if that meant gaining a higher vantage point to see the tops of benches more clearly.) Similar to before, the 50 mm images averaged 5 seconds per station, the 100 mm lens averaged 30 seconds per station, and the 200 mm lens averaged 1 minute 40 seconds per station. Changing lens averaged 1.5 minutes. The entire exercise with all three lenses only took ten minutes because the camera locations were just 25 m apart. In practice it is unlikely that anybody would use a 100 mm lens to capture images from this distance, let alone a 200 mm lens — the ground pixel sizes at the base of the wall for the 50 mm lens, 100 mm lens, and 200 mm lens were 14 mm, 7 mm, and 3.5 mm, respectively. For geotechnical analysis most customers aim for a ground pixel size of 30–50 mm.

Orientations took longer this time due to the need to determine what was surveyed, using the 28 mm project to locate the points. Normally we aim for about 1/3 pixel accuracy but in this case the natural points (e.g. the centre of a large rock) measured tens of pixels across so the digitising accuracy was much lower than normal.

Although the larger projects took quite a long time, it is worth noting that the vast majority of the “Orientations” time for the larger projects and *all* of the “DTM Generation” time is spent with the computer processing the data unattended, and will be quicker with a more powerful PC; actual *user* time for the 200 mm project was about 15 minutes in total. (The reason why the single 100 mm project consisting of two merged images took longer than the ten projects consisting of a single original image with one merged image is that each of those ten projects only had about 450,000 points in them while the single merged image project had over 3 million, and operations like DTM triangulation take longer when there are more points in the DTM.)

Table 5 Image Processing Times

	50 mm (2 merged images)	100 mm (1 merged image)	100 mm (2 merged images)	200 mm (1 merged image[4])
Number of Images	5	20	20	70
Orientations (m:s)	7:15	11:30	12:00	1:06:40
DTM generation (m:s)	6:10	17:55	26:30	1:01:40
Number of Points	651,670	4,533,454	3,096,500	15,940,960

3 Aerial photography

Although terrestrial photography is the preferred option for most applications, there are instances where photographing the surface from above not only simplifies the processing but gives far more uniform results. Examples of this can include subsidence monitoring and volume calculations of pits, stockpiles, and dumps, which is why photogrammetry using aerial photography has been the mainstay of the conventional mapping industry for decades.

The big problem, of course, is the equipment required to obtain aerial photography. In this section we will look at one particular project where ADAM Technology's UAV (Unmanned Aerial Vehicle) technology was used to capture images of stockpiles for volume calculation purposes (Figure 10); the approach applies equally well to subsidence monitoring and volume calculations of other assets, and because the camera can be mounted obliquely, it can also be used for pit wall mapping when large areas need to be covered quickly and suitable vantage points are hard to find.

The UAV itself is an electric model with onboard GPS and magnetometers for navigation. It can take off and navigate a set of waypoints fully autonomously, and land semi-autonomously. (GPS is not accurate enough for it to know when it has reached the ground so the operator brings it down slowly by holding down one button and kills the engine by pressing another once it has landed.) It has a flying time of 15–20 minutes; a two-stroke petrol model with the same basic airframe can fly for around 1.5 hours.

The processing times for this example were measured on a 3.5 GHz Intel Core 2 Quad CPU. These times therefore represent those obtainable from a high-end desktop PC.

3.1 Stockpile volume measurement

To test the feasibility of using UAVs with 3DM Analyst Mine Mapping Suite to calculate stockpile volumes, ADAM Technology conducted an onsite trial at a mine that had crushed ore stockpiles on a 330 m × 260 m pad.

The total time in the field to set up the UAV, run through the pre-flight checklist, and pack up again afterwards, was about 30 minutes, with two people required. (One person acts as the Ground Controller, uploading flight plans and monitoring the UAV's status; the other is the Safety Pilot, whose job is to take manual control of the UAV in an emergency and bring it down safely.)

Nine control points were placed around the stockpiles and surveyed beforehand using GPS.

The flying itself took 7 minutes in total, and the actual aerial photography time was only 3 minutes. (The time required can be seen from Figure 14 because the images were captured on a five-second interval and each image shows up as a camera.)

Processing the images on the quad-core PC took less than 30 minutes in total; 15 million points were generated with a height accuracy of 20 mm and a clear view of the entire stockpile surface, as expected.

---

[4] Actually, two merged images were created but they were both from the same camera station. The reason was that a single merged image was over 320 megapixels and required too much RAM to generate. One of the two merged images was then matched with each original image from the other station so this project is equivalent to the "1 merged image" projects.

Figure 10 ADAM Technology's UAV, the SR20, weighs about 7 kg empty and about 11 kg with camera and batteries. The largest UAV in the range, the SR200, stands about twice as tall and weighs 25 kg.

Figure 11 Colourised point cloud of the stockpiles from above. Average point density is 175 points per square metre, with a height accuracy of 20 mm. Flying height was 120 m. Flying higher would reduce the density and accuracy of the data but also the number of images required and hence the processing time. This particular UAV is capable of covering an area 4–5 times as large as this in a single flight; landing, changing batteries, and taking off again takes about 10 minutes.

Figure 12 The same point cloud from the side.

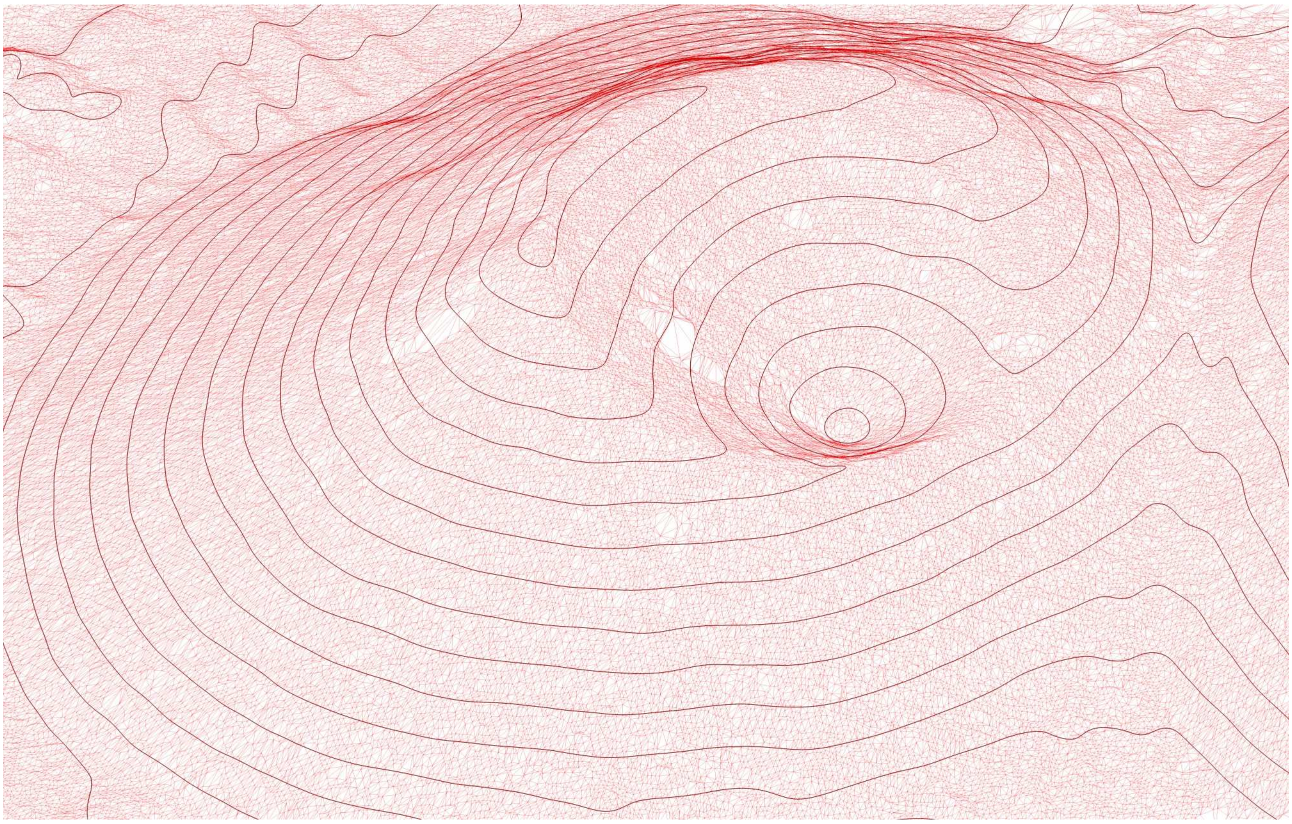

Figure 13 3D surface model of just one of the stockpiles, together with 1 m contours generated by 3DM Analyst. Note the near-uniform point density, possible because of the camera's vantage point looking straight down from above. Terrestrial images and scans would have a much higher point density closer to the camera/scanner but much lower density further away, and would be unable to model the inverted cone at the top of the stockpile.

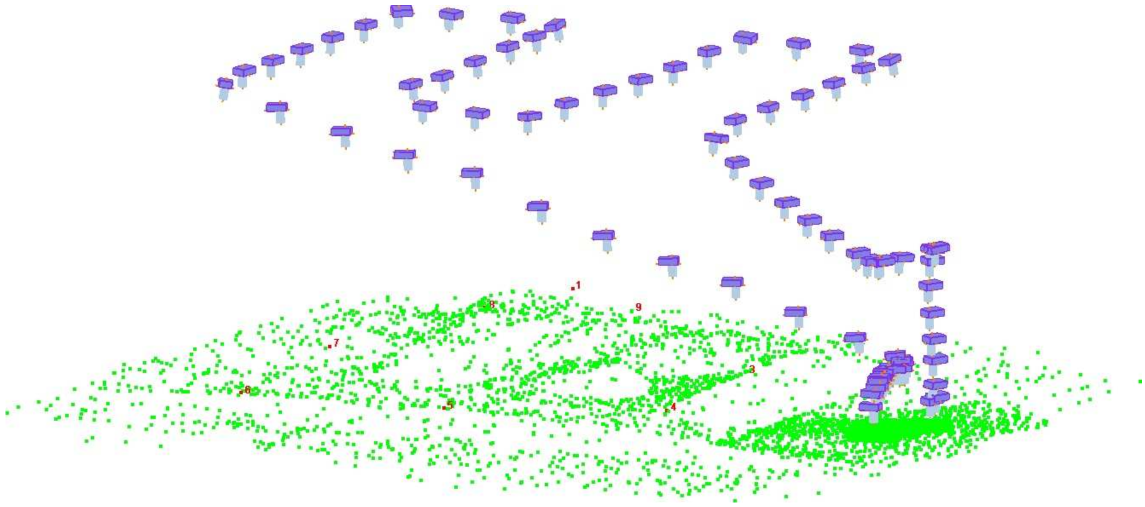

Figure 14 UAV flight path. Each camera represents an image, captured at a five-second interval. The nine control points used are shown in red. Although the flight computer can report the GPS locations of each image, these are generally not accurate enough to be used instead of control points unless accuracy requirements are very low; in this case, all camera locations were determined automatically by the software.

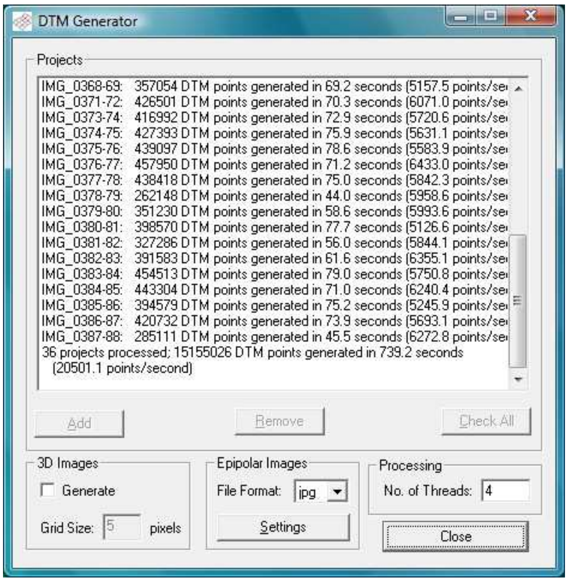

```
IMG_0368-69:   357054 DTM points generated in 69.2 seconds (5157.5 points/se
IMG_0371-72:   426501 DTM points generated in 70.3 seconds (6071.0 points/se
IMG_0373-74:   416992 DTM points generated in 72.9 seconds (5720.6 points/se
IMG_0374-75:   427393 DTM points generated in 75.9 seconds (5631.1 points/se
IMG_0375-76:   439097 DTM points generated in 78.6 seconds (5583.9 points/se
IMG_0376-77:   457950 DTM points generated in 71.2 seconds (6433.0 points/se
IMG_0377-78:   438418 DTM points generated in 75.0 seconds (5842.3 points/se
IMG_0378-79:   262148 DTM points generated in 44.0 seconds (5958.6 points/se
IMG_0379-80:   351230 DTM points generated in 58.6 seconds (5993.6 points/se
IMG_0380-81:   398570 DTM points generated in 77.7 seconds (5126.6 points/se
IMG_0381-82:   327286 DTM points generated in 56.0 seconds (5844.1 points/se
IMG_0382-83:   391583 DTM points generated in 61.6 seconds (6355.1 points/se
IMG_0383-84:   454513 DTM points generated in 79.0 seconds (5750.8 points/se
IMG_0384-85:   443304 DTM points generated in 71.0 seconds (6240.4 points/se
IMG_0385-86:   394579 DTM points generated in 75.2 seconds (5245.9 points/se
IMG_0386-87:   420732 DTM points generated in 73.9 seconds (5693.1 points/se
IMG_0387-88:   285111 DTM points generated in 45.5 seconds (6272.8 points/se
36 projects processed; 15155026 DTM points generated in 739.2 seconds
  (20501.1 points/second)
```

☐ Generate

Grid Size: 5 pixels

File Format: jpg

No. of Threads: 4

Figure 15 DTM Generator processing the models. Four projects were processed at a time.

Conclusions

The first step when using photogrammetry for geotechnical engineering is acquiring the images and georeferencing them. In this paper we have covered the use of image fans to quickly acquire imagery of large areas from longer ranges, and strips of images captured from moving vehicles (both terrestrial and aerial) where speed is of the essence or a suitable vantage point would otherwise be unavailable.

We have also looked at various georeferencing options, from using specially-marked control points surveyed by GPS or a total station, to using natural points or existing features picked up by a reflectorless total station, to surveying three or more camera stations and using no control points at all.

Together these give a range of straightforward options for capturing a scene and successfully georeferencing it, allowing users to focus on the more important problem of actually mapping and analysing the structure of the rock mass.